

Amendments to the Specification:

Please replace the paragraph starting at **line 12 on page 2** with the following amended paragraph:

Buffers, for example in the form of First-In-First-Out (FIFO) queues, are typically used to decouple a source and destination of instructions and/or data, so as to compensate for differences in the delivery rate of a source and the consumption rate of a destination. Within the context of processors (e.g., microprocessors), such buffers are employed at multiple locations between functional units within such processors. Buffers are particularly useful where the output rate of a source functional unit is potentially higher than the consumption rate of a destination functional unit as a result of differences in the output bandwidth and input bandwidth of the source and destination functional units ~~respectfully respectively~~. Further, buffers may be usefully deployed at the point where a source functional unit is clocked at a different speed to a destination functional unit.

Please replace the paragraph starting at **line 24 on page 2** with the following amended paragraph:

Multi-threaded processor design has recently been considered as an increasingly attractive option for increasing the performance of processors. Multithreading within a processor, *inter alia*, provides the potential for more effective utilization of various processor resources, and particularly for more effective utilization of the execution logic within a processor. Specifically, by feeding multiple threads to the execution logic of a processor, clock cycles that would otherwise have been idle due to a stall or other delay in the processing of a particular thread may be utilized to service a further thread. A stall in the processing of a particular thread may result from a number of occurrences within a processor pipeline. For example, a cache miss or a branch ~~misprediction~~ misprediction (i.e., a long-latency operation) for an instruction included within a thread typically results in the processing of the relevant thread stalling. The negative effect of long-latency operations on execution logic efficiencies is exacerbated by the recent increases in execution logic throughput that have outstripped advances in memory access and retrieval rates.